

TAB A

SUBSTITUTE SPECIFICATION

UNITED STATES PATENT APPLICATION

of

DOUGLAS DILLON

for

NETWORK SYSTEM WITH TCP/IP PROTOCOL SPOOFING

This application is a division of Application No. 09/596,603 filed June 19, 2000, which is a division of Application No. 09/204,436 filed December 3, 1998, Patent No. 6,161,141, which is a division of Application No. 08/901,152 filed July 28, 1997, Patent No. 5,995,725, which is a continuation of Application No. 08/257,670 filed June 8, 1994, now abandoned.

Background of The Invention

This application relates to a computer network and, more specifically, to a method and apparatus for allowing both high-speed and regular-speed access to a computer network.

The Internet is an example of a TCP/IP network. The Internet has over 10 million users. Conventionally, access to the Internet is achieved using a slow, inexpensive method, such as a terrestrial dial-up modem using a protocol such as SLIP (Serial Line IP), PPP, or by using a fast, more expensive method, such as a switched 56 Kbps, frame relay, ISDN (Integrated Services Digital Network), or T1.

Users generally want to receive (download) large amounts of data from networks such as the Internet. Thus, it is desirable to have a one-way link that is used only for downloading information from the network. A typical user will receive much more data from the network than he sends. Thus, it is desirable that the one-way link be able to carry large amounts of data very quickly. What is needed is a high bandwidth one-way link that is used only for downloading information, while using a slower one-way link to send data into the network.

Currently, not all users have access to high speed links to networks. Because it will take a long time to connect all users to networks such as the Internet via physical high-speed lines, such as fiber optics lines, it is desirable to implement some type of high-speed line that uses the existing infrastructure.

5 Certain types of fast network links have long propagation delays. For example, a link may be transmitting information at 10 Mbps, but it may take hundreds of milliseconds for a given piece of information to travel between a source and a destination on the network. In addition, for even fast low-density links, a slow speed return-link may increase the round trip propagation time, and thus limit throughput. The
10 TCP/IP protocol, as commonly implemented, is not designed to operate over fast links with long propagation delays. Thus, it is desirable to take the propagation delay into account when sending information over such a link.

Summary of the Invention

15 The present invention overcomes the problems and disadvantages of the prior art by allowing a user to download data using a fast one-way satellite link, while using a conventional low-speed Internet connection for data being sent into the network. The invention uses a "spoofing" technique to solve the problem of the long propagation delays inherent in satellite communication.

20 In accordance with the purpose of the invention, as embodied and broadly described herein, the invention is a network system that forms a part of a network, comprising: a source computer, having a link to the network; a destination computer, having a link to the network; a satellite interface between the source computer and the

destination computer, wherein information passes from the source computer to the destination computer; means in the destination computer for requesting information from the source computer over the network; means for receiving an information packet sent from the source computer in response to the request and for sending the
5 information packet to the destination computer over the satellite interface; and means for sending an ACK message to the source computer in response to receipt of the information packet, wherein the ACK message appears to the source computer to have come from the destination computer.

In further accordance with the purpose of the invention, as embodied and
10 broadly described herein, the invention is a gateway in a network system that forms a part of a TCP/IP network, wherein the network includes a source computer having a link to the TCP/IP network and a link to a high speed satellite interface, and a destination computer having a link to the TCP/IP network and a link to the high speed satellite interface, the gateway comprising: means for receiving an information packet sent from
15 the source computer and for sending the information packet to the destination computer over the satellite interface; and means for sending an ACK message to the source computer in response to receipt of the information packet, wherein the ACK message appears to the source computer to have come from the destination computer.

Objects and advantages of the invention will be set forth in part in the description
20 which follows and in part will be obvious from the description or may be learned by practice of the invention. The objects and advantages of the invention will be realized and attained by means of the elements and combinations particularly pointed out in the appended claims.

Brief Description of The Drawings

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate several embodiments of the invention and, together with the description, serve to explain the principles of the invention.

5 Fig. 1 is a hardware block diagram of a preferred embodiment of the invention;

Fig. 2 is a diagram of a portion of a hybrid terminal of Fig. 1;

Fig. 3 is a diagram showing an IP packet format;

Fig. 4 is a diagram showing a plurality of packet formats, including an Ethernet packet format;

10 Fig. 5 is a diagram showing a tunneling packet format;

Fig. 6 is a diagram of steps performed by the hybrid terminal of Fig. 1;

Fig. 7 is a diagram showing an example of partial data in a tunneling packet;

Fig. 8 is a flowchart of steps performed by the hybrid terminal of Fig. 1;

Fig. 9 is a diagram of steps performed by a hybrid gateway of Fig. 1;

15 Fig. 10 is a diagram showing a format of packets sent to a satellite gateway of Fig. 1;

Fig. 11 is a diagram showing a TCP packet format;

Fig. 12 is a ladder diagram showing packets sent from an application server to the hybrid gateway and from the hybrid gateway to the hybrid terminal over a satellite link; and

20

Figs. 13(a) through 13(e) are flowcharts of steps performed by the hybrid gateway of Fig. 1.

Figs. 14 and 15 are figures from the Phase A Data Sheet incorporated herein.

Detailed Description of The Preferred Embodiments

Reference will now be made in detail to the preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

a. General Overview

A preferred embodiment of the present invention uses satellite technology to implement a high-speed one way link between a user's computer and a TCP/IP network, such as the Internet or a private TCP/IP network. This high-speed link is used to download data from the network. The user's computer also has a conventional TCP/IP link for sending data to the network. The invention can use various forms of high-speed, one-way links, such as satellites, and cable television lines. The invention can use various forms of low-speed networks, such as TCP/IP networks, dialup telephones, ISDN D-channel, CPDP, and low-speed satellite paths.

The described embodiment of the present invention uses satellites to provide a high-speed one-way link. Satellites can cover large geographical areas and are insensitive to the distance between a transmitter and a receiver. In addition, satellites are very efficient at point-to-point and broadcast applications, and are resilient and resistant to man-made disasters. Two-way satellites are expensive to use, however, because of the costs involved in purchasing and installing satellite earth station hardware. In the past, these costs have placed satellite communications outside the reach of the consumer.

The present invention allows a personal computer to receive downloaded information from the network via a satellite at a very practical cost. In the present invention, the cost of satellite communications is reduced because a one-way satellite link is used. Receive-only earth station equipment is cheaper to manufacture because it requires less electronics than send/receive antennae.

As is well-known in the art, communication over the Internet and similar TCP/IP networks is achieved through a group (suite) of protocols called Transmission Control Protocol/Internet Protocol (TCP/IP). The TCP/IP protocol is described in the book "Internetworking With TCP/IP, Vol I" by Douglas Comer, published by Prentice-Hall, Inc., of Englewood Cliffs, N.J., 1991, which is incorporated by reference.

b. Hybrid TCP/IP Access

Fig. 1 is a hardware block diagram of a preferred embodiment of the invention. Fig. 1 includes five subsystems: a hybrid terminal 110, a SLIP provider (Internet connection) 130, an application server 140, a hybrid gateway 150, and a satellite gateway 160. Hybrid terminal 110 is connected to a modem 190, e.g., a 9600 baud modem, which connects to SLIP provider 130 through a telephone line 192. A satellite transmitter 170, a satellite 175, and a satellite receiver 180 provide a fast, one-way link for transferring data from satellite gateway 160 to hybrid terminal 110. Each of SLIP provider 130, application server 140, and hybrid gateway 150 are connected to the Internet 128. As is well-known in the art, the Internet 128 is a "network of networks" and can be visually depicted only in general terms, as seen in Fig. 1.

Each of hybrid terminal 110, SLIP provider 130, application server 140, hybrid gateway 150 and satellite gateway 160 includes a processor (not shown) that executes

instructions stored in a memory (not shown). Other parts of the invention also include processors that are not discussed herein, such as I/O processors, etc. Preferably, hybrid terminal 110, hybrid gateway 150, and satellite gateway 160 are implemented as personal computers including an 80386/80486 based personal computer operating at
5 least 33 MHz, but these elements can be implemented using any data processing system capable of performing the functions described herein. In the described embodiment, SLIP provider 130 is a conventional SLIP provider and application server 140 is any application server that can connect to the Internet 128 via TCP/IP.

As shown in Fig. 1, hybrid terminal 110 preferably includes application software
10 112, driver software 114, a serial port 122 for connecting hybrid terminal 110 to modem 190, and satellite interface hardware 120 for connecting hybrid terminal 110 to satellite receiver 180.

Fig. 2 shows a relationship between software in application 112, software in driver 114, serial port 122, and satellite interface 120. Application software 112 includes
15 TCP/IP software, such as SuperTCP, manufactured by Frontier, Inc., Chameleon, manufactured by Netmanager, and IRNSS, manufactured by Spry, Inc. The described embodiment preferably operates with the SuperTCP TCP/IP package and, thus, uses a standard interface 212 between the TCP/IP software 210 and driver 114. Examples of standard interface 212 between TCP/IP software 210 and driver 114 includes the
20 Crynson-Clark Packet Driver Specification and the 3Com/Microsoft Network Driver Interface Specification (NDIS). Other embodiments use other standard or non-standard interfaces between TCP/IP software 210 and driver 114.

As shown in Fig. 2, application software preferably 112 also includes well-known Internet utilities, such as FTP 230, and well-known user interfaces, such as Mosaic and Gopher (shown). Application software 112 can also include other utilities, e.g., News and Archie (not shown).

5 The following paragraphs describe how a request from hybrid terminal 110 is carried through the Internet 128 to application server 140 and how a response of application server 140 is carried back to the user at hybrid terminal 110 via the satellite link. The operation of each subsystem will be described below in detail in separate sections.

10 In the present invention, hybrid terminal 110 is given two IP addresses. One IP packet address corresponds to SLIP provider 130 and is assigned by a SLIP service provider. The other IP address corresponds to satellite interface 120 and is assigned by a hybrid service provider. IP addresses are assigned by the SLIP and satellite network managers and loaded into hybrid terminal 110 as part of an installation
15 configuration of the hybrid terminal's hardware and software. These two IP addresses correspond to completely different physical networks. SLIP provider 130 does not "know" anything about the satellite IP address or even whether the user is using the satellite service. If a host somewhere in the Internet is trying to deliver a packet to the satellite interface IP address by using the Internet routing scheme of routers, gateways,
20 and ARPs (Address Resolution protocol), the only way that the packet can reach the satellite IP interface is to traverse the satellite by being routed through satellite gateway 160.

The following example assumes that a user at hybrid terminal 110 desires to send a request to a remote machine, such as application server 140, that is running FTP (File Transfer protocol) server software. The FTP software running on application server 140 receives file transfer requests and responds to them in an appropriate fashion.

Fig. 3 shows the contents of a source field (SA) and of a destination field (DA) of packets sent between the elements of Fig. 1. A request for a file and a response of a file sent from application server 140 to hybrid terminal 110 takes the following path.

- 1) Within hybrid terminal 110, FTP client software 230 generates a request and passes it to TCP/IP software 210. TCP/IP software 210 places the request in a TCP packet (see Fig. 11). Next, the TCP packet is placed in an IP packet, having a format shown in Fig. 3. TCP/IP software 210 places the IP packet in an Ethernet packet, as shown in Fig. 4, and passes the Ethernet packet to driver 114. This packet has a source IP address corresponding to satellite interface 120 and a destination IP address of application server 140.
- 2) In driver 114, the Ethernet header and checksum are stripped off the packet and the IP packet is encapsulated, or "tunneled," inside of another IP packet and sent over serial port 122 to SLIP provider 130. Fig. 5 shows a format of a tunnelled packet. Fig. 7 shows an example of a tunnelled packet. The encapsulation adds a new IP header 530 in front of the original packet 540 with a source address corresponding to SLIP

provider 130 and a destination address corresponding to hybrid gateway 150.

3) SLIP provider 130 receives the IP packet, analyzes the tunneling header and, thinking it is destined for hybrid gateway 150, uses standard Internet routing to send the packet to hybrid gateway 150.

4) When hybrid gateway 150 receives the packet, it strips off the tunneling header, revealing the true header with application server 140 as the destination. The packet is then sent back out into the Internet 128.

5) Internet routing takes the packet to application server 140, which replies with the requested file and addresses the reply to the request's source IP address, i.e., the IP address of the hybrid terminal's satellite interface 120.

6) In order to find the hybrid terminal's satellite interface 120, the Internet routing protocol will send the packet to the subnet containing a router/gateway connected to hybrid gateway 150. When a router on the same physical network as satellite gateway 160 and hybrid gateway 150 sends out an ARP for the IP address of satellite interface 120 (to find a physical address of satellite interface 120), hybrid gateway 150 responds and says "send it to me." Thus, application server 140 and the rest of the Internet 128 think that packets sent to hybrid gateway 150 will reach the hybrid terminal's satellite interface.

7) Once hybrid gateway 150 receives a reply packet from application server 140, it sends it to satellite gateway 160. In the described embodiment, hybrid gateway 150 encapsulates the packet in a special packet format

that is used over the satellite link and uses the satellite interface IP address to uniquely identify the satellite packet's destination. Then hybrid gateway 150 sends the packet over the Ethernet to satellite gateway 160.

- 5 8) Satellite gateway 160 broadcasts over the satellite link any packets it receives from hybrid gateway 150.
- 9) Driver 114 in hybrid terminal 110 that services satellite interface 120 scans all packets broadcast over satellite transmitter 170 looking for its satellite interface IP address in the header. Once it identifies one, it captures it, strips off the satellite header revealing the reply IP packet, and
10 sends it to driver 114.

Thus, IP packets sent into Internet 128 are carried by the SLIP connection, while IP packets from the Internet 128 are carried by the satellite link. The following paragraphs describe the operation of each subsystem in more detail.

1. The Hybrid Terminal

15 Hybrid terminal 110 is the terminal with which the user interacts. Thus, hybrid terminal 110 includes a user interface device (not shown) such as a mouse, keyboard, etc. As shown in Fig. 1, hybrid terminal 110 includes one or more application programs 112 (including TCP/IP software 210), and driver software 114, which communicates with SLIP provider 130 through a serial port 122 and modem 190, using a driver portion
20 118, and which communicates with satellite receiver 180 through a satellite interface 120, using a driver portion 116.

To TCP/IP software 210, driver 114 appears to be an Ethernet card, although driver 114 is actually connected to satellite receiver 180 (via satellite interface 120) and

to SLIP provider 130 (via serial line 122 and modem 190). Thus, TCP/IP software 210 believes that it is communicating with a single physical network, when it is, in reality, communicating with two physical networks (the SLIP dial-up network and a satellite network). Ethernet is a packet switching protocol standardized by Xerox Corporation, Intel Corporation, and Digital Equipment Corporation, which is described in "The Ethernet: A Local Area Network Data Link Layer and Physical Layer Specification," September 1980, which is available from any of these three companies, and which is incorporated by reference.

Fig. 6 is a diagram of steps performed by driver 114 of hybrid terminal 110 of Fig.

1. As shown in Fig. 6, driver 114 receives packets of data from TCP/IP software 210 and passes them to SLIP provider 130 via serial port 122 and modem 190. A packet sent by application server 140 is received through satellite receiver 180, passed through the satellite interface 120, passed to the satellite driver 220, and passed to driver 114, which passes the received packet to TCP/IP software 210.

The following paragraphs discuss two basic functions performed by driver 114 (tunneling and ARP handling) and discuss various implementation details for the preferred embodiment.

A. "Tunneling"

As discussed above, hybrid terminal 110 has two IP addresses associated with it: one for SLIP provider 130 and one for the satellite interface 120. Packets containing requests are sent from hybrid terminal 110 to application server 140 via the Internet 128, while packets containing a reply are sent back via the satellite link. Tunneling is

the method by which application server 140 is "fooled" into sending a reply to a different IP address (satellite interface 120) than that of the sender (serial port 122).

A packet received by driver 114 from the TCP/IP software 210 has a source address of satellite gateway 160 and a destination address of application server 140.

5 As shown in step 610 of Fig. 6, driver 114 removes the Ethernet header and checksum and encapsulates the IP header into an IP tunneling header having a source address of SLIP provider 130 and a destination address of hybrid gateway 150 (see Fig. 7). As described above, at hybrid gateway 150, the tunneling header is removed and the packet is sent back into the Internet 128 to be sent to application server 140.

10 When forming a tunneling header, driver 114 copies all the values from the old header into the new one with the following exceptions. The source and destination addresses of the tunneling header change, as described above. In addition, a total packet length field 510 is changed to contain the contents of length field 310 plus the length of the tunneling header. Lastly, the driver 114 recalculates checksum 520 of the
15 tunneling header because some of the fields have changed.

B. ARP Handling

ARP (Address Resolution Protocol) is used by TCP/IP to dynamically bind a physical address, such as an Ethernet address, to an IP address. When TCP/IP finds an IP address for which it does not know a physical address, TCP/IP broadcasts an
20 ARP packet to all nodes, expecting a response that tells TCP/IP what physical address corresponds to the IP address.

During initialization, driver 114 declares to TCP/IP software 210 that driver 114 is an Ethernet card to ensure that the packets that TCP/IP package sends are Ethernet

packets and that the TCP/IP package will be prepared to receive packets at a high-rate of speed. As shown in step 620 of Fig. 6, when driver 114 detects that TCP/IP has sent an ARP packet, driver 114 creates a physical address and sends a reply packet to TCP/IP software 210. The contents of the physical address are irrelevant, because
5 driver 114 strips off the Ethernet header on packets from TCP/IP before the packets are sent to SLIP provider 130.

C. Other Functions

As shown in step 630 of Fig. 6, packets received by driver 114 from satellite receiver 180 (via satellite driver 114) are merely passed to TCP/IP software 210. The
10 following paragraphs discuss implementation details for the described embodiment.

In a preferred embodiment, TCP/IP software 210 (e.g., Frontier's SuperTCP) sends an ACK (acknowledge) for every packet it receives, even though this action is not required by the TCP/IP protocol. In this situation, many packets compete for the slow link to SLIP provider 130. In TCP/IP, the ACK scheme is cumulative. This means that
15 when a transmitter receives an ACK stating that the receiver has received a packet with sequence number N, then the receiver has received all packets with sequence numbers up to N as well, and there is no reason why every packet needs to be ACK'ed.

Fig. 8 is a flowchart of steps performed in a preferred embodiment by driver 114 of hybrid terminal 110. Fig. 11 is a diagram showing preferred a TCP packet format.
20 Fig. 11 includes a sequence number field 1102, an acknowledgment (ACK) number field 1104, and a checksum field 1106. In step 810 of Fig. 8, driver 114 receives an ACK packet with sequence number N from TCP/IP software 210. The packet is queued along with other packets waiting to be sent to SLIP provider 130. In step 820 driver 114

checks to determine whether there is a "run" of sequential packets waiting to be sent. If so, in step 830, driver 114 deletes ACK packets for the same TCP connection that have sequence numbers in the run from the queue and sends an ACK only for the highest sequence number in the run. This action alleviates the bottleneck caused by the relatively slow modem speeds.

Serial port 122 provides a physical connection to modem 190 and, through it, to the terrestrial network via a SLIP protocol as described below in connection with SLIP provider 130. Serial data is sent and received through an RS-232 port connector by a UART (Universal Asynchronous Receiver Transmitter), such as a U8250, which has a one byte buffer and is manufactured by National Semiconductor, or a U16550, which has a 16 byte buffer and is also manufactured by National Semiconductor.

The invention preferably operates under the DOS operating system and Windows, but also can operate under other operating systems.

Satellite driver software 220 receives packets from satellite 180, and passes them to driver 114 using a DOS call. Thus, the two physical links are combined within driver 114 and the existence of two physical links is transparent to TCP/IP software 210. Satellite driver 220 scans all packets transmitted over the satellite channel for a packet with a header corresponding to the IP address of the satellite interface 122, performs some error detection and correction on the packet, buffers the received packet, and passes the packet to driver 114 using a DOS call, e.g., IOCTL-output-cmd(). Driver 114 copies data from satellite driver 220 as quickly as possible and passes it to TCP/IP software 210.

As discussed above, TCP/IP software 210 is fooled into thinking that it is connected to an Ethernet network that can send and receive at 10 Mbps. This concept is helpful on the receive side because data from the satellite is being received at a high rate. On the transmit side, however, modem 190 is not capable of sending at such a high rate. In addition, TCP/IP software 210 sends Ethernet packets to driver 114, i.e., an IP packet is encapsulated into an Ethernet packet. Because SLIP provider 130 expects IP packets, driver 114 must strip the Ethernet header before the packet is sent to SLIP provider 130.

As described above in connection with Fig. 8, driver 114 also includes a transmit and receive queue. As data is received from TCP/IP software 210 and received from the satellite driver 220, it is buffered within the queue. When the queue is full, e.g., when TCP/IP is sending packets faster than modem 190 can send them, driver 114 drops the packets and returns an error so that TCP/IP software 210 will decrease its rate of transmission.

In a first preferred embodiment, a SLIP connection is initiated with an automatic logon procedure. In another preferred embodiment, driver 114 executes instructions to allow a user to perform a SLIP logon manually.

Because TCP/IP software 210 preferably is configured to talk to Ethernet and it is desirable to receive the largest packet size possible, driver 114 configures TCP/IP so that the MTU (Maximum Transmission Unit) of the network is as large as possible, e.g., 1500 bytes. Some SLIP providers 130 have a smaller MTU, e.g., 512 bytes. To handle the disparity in size, driver 114 segments large packets received from TCP/IP software 210 into segments the size of the SLIP MTU. Once a packet is segmented, it is

reassembled in hybrid gateway 150. Only the tunneling header is copied as the header of the segments.

2. The SLIP Provider

SLIP provider 130 performs the function of connecting hybrid terminal 110 to the Internet 128. As described above, other protocols, such as PPP, could also be used to perform the connecting function. SLIP server 130 receives SLIP encoded IP packets from modem 190, uncodes them, and forwards them to hybrid gateway 150 via the Internet 128.

In its most basic form, SLIP provider 130 delimits IP packets by inserting a control character hex 0xC0 between them. To insure that a data byte is not mistaken for the control character, all outgoing data is scanned for instances of the control character, which is replaced by a two character string. The SLIP protocol is described in detail in J. Romkey, "A Nonstandard for Transmission of IP Datagrams over Serial Lines: SLIP," RFC 1055, June 1988, pp. 1-6, which is incorporated by reference.

3. The Application Server

Application server 140 is a computer system running any combination of known application programs available on the Internet using the TCP/IP protocol suite. For example, application server 140 may be transferring files to requesting users via FTP. Although hybrid terminal 110 Actually has two IP addresses (a serial port address and an address for the satellite interface), the software executing on application server 140 thinks that it is receiving requests over the satellite network and sending responses over the satellite network. Hybrid terminal is completely transparent to application server 140.

4. The Hybrid Gateway

Although only one hybrid terminal 110 is shown in Fig. 1, the invention can include a plurality of hybrid terminals 110. Preferably, all packets sent from all hybrid terminals 110 pass through hybrid gateway 150 to get untunnelled. Thus, hybrid gateway 150 is a potential system bottleneck. Because of this potential bottleneck, the functions of hybrid gateway 150 are as simple as possible and are performed as quickly as possible. Hybrid gateway 150 also has good Internet connectivity to minimize the accumulated delay caused by packets waiting to be processed by hybrid gateway 150.

A. Untunnelling

Fig. 9 is a diagram of steps performed by hybrid gateway 150 of Fig. 1. In step 910, hybrid gateway 150 receives a tunneled packet having a format shown in Fig. 5. Hybrid gateway 150 "untunnels" the packet by stripping off the tunneling header and passes the packet back to the Internet 128.

As described above, packets are sometimes broken into segments when they are sent in order to accommodate a small MTU of SLIP provider 130. Packets may also be segmented as they pass through other elements of the Internet 128 having small MTUs. For fragmented packets, only the tunnelled header is copied into the header of each segment. Hybrid gateway 150 stores fragmented packets in a memory (not shown) and reassembles them in order before untunnelling the original packet and passing it to the Internet 128. Preferably, a "time to live" value is assigned to each packet when it is sent by driver 114 and if all segments do not arrive before a time to live timer expires, the packet is discarded.

B. ARP Responding

Preferably, satellite gateway 160 is on a same physical network as hybrid gateway 150. As shown in step 920 of Fig. 9, when a router on the same physical network as satellite gateway 160 and hybrid gateway 150 sends out an ARP for the IP address of satellite interface 120 (to find a physical address of satellite interface 120), hybrid gateway 150 responds and says "send it to me." Hybrid gateway 150 needs to intercept packets intended for satellite interface 120 because it needs to encapsulate packets for satellite gateway 160 as follows.

C. Satellite Packetizing

The following paragraphs describe how packets travel from application server 140 through hybrid gateway 150 and to satellite gateway 160. The following explanation is given by way of example and is not intended to limit the scope of the present invention. As shown in step 930 of Fig. 9, hybrid gateway 150 encapsulates replies from application server 140 into a satellite packet format. Fig. 10 is a diagram showing a format of a satellite packet sent to satellite gateway 160 of Fig. 1. A satellite packet includes the data 1010 of an original IP packet and two headers 1020, 1030 added by hybrid gateway 150.

Satellite gateway 160 expects IP packets to be encapsulated first in a special satellite packet and then within an LLC-1 IEEE 802.2 link level control, type 1 packet.

Satellite header 1020 identifies the downlink and contains a sequence number and the packet length. An LLC-1 header 1030 preferably is used to send the packet to satellite gateway 160, in an Ethernet LAN. Hybrid gateway 150 prepares packets for satellite gateway 160 by appending headers 1020 and 1030 to the front of an IP packet 1010.

The receiver in hybrid terminal 110 does not receive the LLC-1 header 1030. Hybrid terminal 110 identifies packets intended for it by checking a least significant byte in the satellite IP address. Thus, a six byte satellite destination address is determined by reversing an order of bytes of the satellite IP address for hybrid terminal 110 and then padding the rest of the address with zeroes.

5. The Satellite Gateway

Satellite gateway 160 can include any combination of hardware and software that connects satellite transmitter 170 to hybrid gateway 150. Satellite transmitter 170 and satellite receiver 180 can be any combination of hardware and software that allows data to be transmitted by satellite transmitter 170 and received by satellite receiver 180, and to be input to hybrid terminal 110. For example, satellite gateway 160 preferably is a personal computer with a high-speed Ethernet connection to hybrid terminal 110. When satellite gateway 160 receives a packet from hybrid gateway 150, it sends it over the satellite link.

Satellite communication may be effected by, for example, the Personal Earth station manufactured by Hughes Network Systems Inc. In a preferred embodiment, a one-way version of the Personal Earth Station is used. Another embodiment uses a satellite communication system manufactured by Comstream. Yet another embodiment uses a system that allows hybrid terminal 110 to be connected directly to satellite receiver 180 via Hughes Network Systems' DirecPC product. The DirecPC satellite interface card is described in "DirecPC, Phase A Data Sheet," dated June 7, 1993, which is incorporated by reference and by the inclusion of its contents which read as follows:

"DirecPC is a satellite, one-way broadcast network offering three services to the IBM compatible PC:

1. Digital package delivery -- Software, games, multi-media news, electronic documents and any other data in the form of a collection of PC files are made available to the PC on a scheduled or on-demand basis.
2. Data Pipe -- provides multiple independent digital streams to carry video, audio, etc.
3. Hybrid Internet Access -- high-speed, low-cost Internet connection where DirecPC carries packets from the Internet and dial-up modem carries packets into the Internet.

See Fig. 14.

To receive the DirecPC broadcast, a PC is equipped with a PC plug-in card and a 24 inch antenna. DirecPC uses a full Galaxy class Ku-Band transponder to provide an 11 Mbps broadcast channel. DES encryption based conditional access ensures that a receiver PC may only access data it is authorized to receive.

Section 1 PC User Perspective

The PC hardware consists of the DirecPC adapter, an antenna and a TVRO standard coaxial cable. The DirecPC adapter is a 16-bit ISA adapter providing throughput comparable to a 16-bit ISA ethernet adapter.

5 The software appears to the user as a set of Windows applications. The applications:

- assist installation and service registration.
- support package delivery by allowing the user to select packages for reception, be notified when packages are received. The
10 software also supports billing for packages received.
- provide a TCP/IP protocol stack and set of applications for Hybrid Internet access.
- provide a driver DLL on which third party software may layer data pipe applications.

15 The software for a data pipe service is provided by the enterprise providing the service. Communications back to the uplink is required for

billing purposes and also for Hybrid Internet access. These communications take place via the PC's dial-up AT command-set modem.

Section 2 Open Interfaces And APIs

5 The DirecPC architecture is open, allowing content providers complete control over their content and the user interface to their content. DirecPC provides interfaces to content providers at the uplink and Application Programming Interfaces (APIs) on the receiving PC. The specifications and APIs are available on request.

See Fig. 15.

10 Section 3 Content Providers

A content provider is an organization that supplies the data sent over the DirecPC system. A content provider can be categorized as being either a:

1. Package Publisher -- uses the DirecPC system as a means of selling and distributing software packages or data packages where
15 a package consists of a set of PC files.

2. Data Pipe Provider -- uses the DirecPC system as a data pipe transport mechanism. User services (News Feeds, Internet Access, Broadcast Video and Audio, etc.) are layered on top of a datagram transport.

5 DirecPC supports multiple content providers of both kinds.

Section 4 DirecPC Package Distribution

10 The DirecPC system allows data packages to be distributed and purchased. The term "package" refers to any data (including electronic documents, multi-media data, software packages, games, etc.) which can take the form of a group of PC files.

To prepare a package for transmission, a publisher merges the package's files into a single file using the appropriate utility (e.g. PKZIP or ARJ) and loads the package into the uplink using an off-the-shelf file transfer mechanism (e.g. TCP/IP's FTP, floppy-disk, CD-ROM, X-Modem, etc.).

15 Scheduling, pricing and conditional access restrictions can be performed either manually or automatically under publisher control when the package is loaded into the uplink.

DirecPC's conditional access mechanism ensures that a user may only receive authorized packages. As part of initial registration, the user is provided a credit limit. The PC locally maintains a credit account. When the user selects a package for reception, the PC records the transaction and debits the account. A log of all package receptions is maintained on the PC's hard disk and can be browsed by the graphical front-end.

On uplink operator command, when the local credit limit is exceeded or when the user has purchased a certain number of packages, the PC makes a dial-up call to the DirecPC billing service. The call reports the billing information as well as usage information of packages received.

The usage information is used to provide feedback for future scheduling of packages. The reports given to publishers include for each package reception, the name, address etc. of the recipient, the ID of the package and when package delivery took place.

A software package may either be transmitted on a scheduled basis or on-demand. Scheduled transfers are perfect for:

1. Periodical Distribution -- examples include news and weather updates, electronic newspaper, magazine and catalog distribution.

2. Popular Package Delivery -- packages for which there are expected to be multiple recipients. The most popular (or highest profit) packages would be scheduled more frequently to reduce the average time spent waiting, while less popular packages may be scheduled for overnight delivery. Scheduled delivery is lower cost than delivering a package on-request to each buyer. The schedule for individual packages is manually set by hub operators with the submission of the package.

Phase A package delivery allows a single transmission at any given time. The rate of transmission is settable under operator control at speeds up to 2 Mbits/sec. Support for simultaneous transmissions will be provided in a subsequent release of DirecPC software.

A software package may be transmitted on-demand in the gaps between scheduled transmissions. Such a transfer delivers the information more quickly to the requesting PC, but at greater cost as the package is not broadcast. A PC uses its modem to request the package.

DirecPC's low bit error rate and high availability ensure that packages are reliably delivered with one transmission. For even greater reliability, each package may be set to employ one or more of the following methods to ensure fail-safe delivery:

1. Repeated Transmission -- A package may be scheduled to be sent more than once to ensure its delivery. A receiving PC, if any packets are lost on the first transmission, fills in the gaps on subsequent transmissions. This mechanism ensures extremely high probability of delivery without requiring use of a return link.
2. Retransmission requests -- a PC, if it misses parts of a package, may request retransmission of those parts. The missing parts are multi-cast so that parts need only be retransmitted once even though they were missed by multiple PCS. Retransmission requests are most appropriate for scheduled individual package transmissions where the package is scheduled less frequently.
3. Delivery confirmation -- a PC, after successfully receiving and installing a package, may send a confirmation to the hub. These confirmations are tabulated and provided in the form of reports to the publisher. This method is more expensive in that it requires that a delivery confirmation (entailing a separate call) be sent by every receiving PC.

Section 5 Data Pipe Transmission

DirecPC's data pipe services are modelled on Local Area Network multi-cast transmission. The data pipe provider passes 802.2 LLC1 Token-Ring or Ethernet multi-cast packets to the uplink. This allows off-the-shelf bridges and routers to be used to support a terrestrial backhaul. It also allows some LAN based applications to operate across the spacelink with little or no modification. The uplink relays these packets across the spacelink. The DirecPC driver passes received packets to the applications. To prevent unauthorized access, each multi-cast address is encrypted under a different key. The DirecPC device driver API allows applications to designate which multi-cast addresses are of interest. Hardware filtering in the DirecPC adapter allows the reception of any 100 different multi-cast addresses.

DirecPC network management allocates to each service provider:

1. a Committed Information Rate (CIR) -- a fraction of broadcast channel bandwidth which is guaranteed to the data pipe provider, and
2. one or more multi-cast 48 bit addresses -- each address operates as a separate data stream multiplexed on the one broadcast channel.

Section 6 Hybrid Internet Access

Hybrid Internet access allows a PC high-speed (over 100 Kbps) access to the Internet. An HNS (Hughes Network Systems) provided NDIS device driver operates with an off-the-shelf TCP/IP package. Reception from the Internet takes place via DirecPC. Transmission into the Internet takes place via a dial-up SLIP connection into the uplink. Hybrid Internet Access allows operation of all the standard Internet applications including SMTP EMAIL, NNTP Usenet News, FTP, GOPHER and Mosaic. As part of initial registration, each receiving PC is provided a permanently assigned IP address.

Hybrid Internet Access is the result of joint development by HNS and the University of Maryland funded in part by a MIPs grant. Continuing development will increase performance and allow receive-only reception of Usenet News.

Section 7 Performance Specifications

Averaged across a whole year, each DirecPC receiver should be expected to have a BER less than $10E-10$ more than 99.5% of the time where a single bit error causes the loss of an entire packet.

Section 8 User Characteristics

The receiver (antenna, cabling and PC plug-in card) is intended to be self-installable by consumers and small business. In cases where self-installation is not desirable, the DirecPC adapter will be installed by the customer and the antenna and cable will be installed by the HNS VSAT installers. The customer uses diagnostic software provided with the adapter to ensure that the PC as a whole is ready for the antenna to be installed.

Maintenance will be performed either by the user swapping components (DirecPC adapter, LNB, etc. with telephone support). HNS's nationwide VSAT field-service network may also be contracted for."

At the downlink, satellite receiver 180 includes a 0.6 meter receive-only antenna receiving HDLC encapsulated LAN packets. Satellite interface 120 includes rate 2/3 Viterbi/Reed-Soloman concatenated forward error correction.

Although only one hybrid terminal 110 and one application server 140 are shown in Fig. 1, the invention can include a plurality of hybrid terminals 110 and/or a plurality of application servers 140. Preferably, all packets sent from all application servers 140 to a hybrid interface 110 pass through satellite gateway 160. Thus, satellite gateway 160 is a potential system bottleneck. Because of this potential bottleneck, the functions of satellite gateway 160 are as simple as possible and are performed as quickly as possible.

c. Protocol Spoofing

TCP/IP protocol specifies that only a predetermined number of packets can be outstanding during transmission, i.e., that only a limited number of packets can be sent before an ACK (acknowledgment) is received. The high bandwidth and long delays incurred in sending packets to an orbiting satellite and back means that at any given time, a large number of packets are "in the pipe" between transmitter and receiver.

When using conventional TCP/IP protocol, application server 140 sends a predetermined number of packets in accordance with a predetermined window size, and then waits to receive ACKs over the modem link before sending additional packets. The purpose of windowing is to limit a number of packets that must be re-sent if no ACK is received and to provide flow control, e.g., to prevent sending packets faster than they can be received. The packets that have not been ACK'ed are stored in a memory so that they can be re-sent if no ACK is received.

In a preferred embodiment of the present invention, hybrid gateway 150 "spoofs" application server 140 to improve the throughput over the satellite link. Specifically, hybrid gateway 150 sends an ACK to application server 140, even though a corresponding packet may not have been received by hybrid terminal 110 via the satellite at the time.

Fig. 12 is a ladder diagram showing packets sent from application server 140 to hybrid gateway 150 and from hybrid gateway to hybrid terminal 110 through the satellite link. Fig. 12 is not drawn to scale. In Fig. 12, application server 140 sends a message #1 to hybrid gateway 150. The propagation time for this transmission is relatively short. Hybrid gateway 150 immediately creates an ACK packet and sends it to application

server 140. Hybrid gateway 150 also sends packet #1 to hybrid terminal 110 through the satellite link. This transmission has a long propagation delay. When hybrid terminal 110 receives the packet, it sends an ACK #1 back to hybrid gateway 150 (e.g., using the tunneling mechanism described above). In a system that does not use tunneling,
5 hybrid gateway 150 needs to intercept the ACK packets from hybrid terminal 110.

Figs. 13(a) through 13(e) are flowcharts of steps performed by hybrid gateway 150 of Fig. 1 during protocol spoofing. In step 1302 of Fig. 13(a), hybrid gateway 150 receives a packet from application server 140 indicating that a new connection is being formed between application server 140 and hybrid terminal 110. In step 1304, hybrid
10 gateway 150 sets up a queue or similar data structure in memory to save un-ACK'ed packets for the new connection. Fig. 13(b) shows corresponding steps performed by hybrid gateway 150 when the connection is closed. Hybrid gateway 150 receives a packet indicating the closure in step 1306 and deletes the queue and saved values for the connection in step 1308.

15 In step 1310 of Fig. 13(c), hybrid gateway 150 fails to receive an ACK for a packet number X from hybrid terminal 110 before an end of a predetermined timeout period. Hybrid gateway 150 maintains a timer for each un-ACK'ed packet. At the end of the predetermined period, hybrid gateway 150 retransmits a packet corresponding to the expired timer. In step 1312, hybrid gateway 150 re-sends packet number X, which it
20 previously saved in the memory queue for this connection (see Fig. 13(d) below).

In step 1314 of Fig. 13(d), hybrid gateway 150 receives a packet from application server 140. In step 1316, hybrid gateway 150 sends the received packet to satellite gateway 160, where it is transmitted over the satellite link, and saves the packet in case

it needs to be retransmitted (see Fig. 13(c)). Hybrid gateway 150 then creates an ACK packet to send to application server 140 in step 1318. The created ACK packet incorporates a format shown in Fig. 11. Hybrid gateway 150 creates an ACK number for field 1104. The ACK number is determined as follows:

5 Hybrid gateway 150 saves the following information for each connection:

1) Send sequence number - a highest in-sequence sequence number of packets sent by application server 140 over the connection.

2) ACK sequence number - the ACK sequence number from the most recent packet sent by hybrid terminal 110 over this connection.

10 3) ACK window size - the window size from the most recent packet from hybrid terminal 110 over this connection.

4) ACK number - the ACK sequence number that is relayed to application server 140. The ACK number is set to:

15 $\text{minimum}(\text{send sequence number}, \text{ACK sequence number} + \text{spoofed window size} - \text{ACK window size})$.

5) spoofed window size - predetermined maximum number window size to be allowed on this connection.

When hybrid gateway 150 inserts the ACK number in the packet, it also calculates the packet's checksum 1106.

20 In step 1320 of Fig. 13(e), hybrid gateway 150 receives an ACK packet over the modem link from hybrid terminal 110. In step 1322, hybrid gateway 150 removes from the queue the packet for which the ACK was received. Because an ACK was received, the packet does not need to be re-sent. In the TCP/IP protocol, a packet containing an

ACK may or may not contain data. Hybrid gateway 150 edits the received packet to replace the packet's ACK number 1104 with a "spoofed" ACK number in step 1326. The spoofed ACK number is determined in the same way as the ACK number in step 1318 of Fig. 13(d). When hybrid gateway 150 substitutes the spoofed ACK number 1104 in the packet, it also recalculates the packet's checksum 1106 in step 1326.

In step 1328, hybrid gateway 150 forwards the received ACK packet to application server 140. Application server 140 may simply disregard the packet if it contains an ACK and no data. In another embodiment, hybrid gateway 150 simply discards a packet received from hybrid terminal 110 that contains an ACK, but no data.

If the connection goes down, either explicitly or after a predetermined period of time, hybrid gateway 150 deletes the saved packets for the connection.

d. Summary

In summary, the present invention allows a personal computer to send messages into the Internet using a conventional dial-up link and to download data from the Internet using a high-speed one-way satellite link. In a preferred embodiment, the invention uses a conventional SLIP provider to connect to the Internet and uses a commercial software TCP/IP package that has a standard driver interface. A spoofing protocol compensates for the long propagation delays inherent to satellite communication.

Other embodiments will be apparent to those skilled in the art from consideration of the specification and practice of the invention disclosed herein. It is intended that the specification and examples be considered as exemplary only, with a true scope of the invention being indicated by the following claims.